Exhibit 1

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

NXP USA, INC. and QUALCOMM INCORPORATED,

Petitioners,

v.

REDSTONE LOGICS LLC,

Patent Owner.

Case: IPR2025-00485
U.S. Patent No. 8,549,339

DECLARATION OF MURALI ANNAVARAM, Ph.D.

TABLE OF CONTENTS

I. INTRODUCTION

1. My name is Dr. Murali Annavaram, Ph.D. I have been retained by the Petitioners as an independent expert consultant in this inter partes review ("IPR") proceeding before the United States Patent and Trademark Office ("PTO") regarding US Patent No. 8,549,339 ("the '339 Patent").

II. PROFESSIONAL BACKGROUND

2. All of my opinions stated in this Declaration are based on my own personal knowledge and professional judgment. In forming my opinions, I have relied on my knowledge and experience in designing, developing, researching, and teaching the technology referenced in this Declaration.

3. I am over 18 years of age and, if I am called upon to do so, I would be competent to testify as to the matters set forth herein. I understand that a copy of my current curriculum vitae, which details my education and professional and academic experience, is being submitted as EX1003. The following provides a brief overview of some of my experience that is relevant to the matters set forth in this Declaration.

4. I am an expert in the field of computer architecture, including power efficient multi-core computer architecture and design. My research focuses on efficient power management of chip multiprocessors (CMPs) and graphics processing units (GPUs), power efficient heterogeneous CMP designs, system-level

5. Element-by-Element Analysis

81. The analysis of each element below incorporates the above descriptions of each reference, the motivation to combine the references, and the resulting modification.

a) Claim 1

[1.pre] A multi-core processor, comprising:

82. In my opinion, White/Talwar renders obvious Limitation [1.pre]. As discussed in Section X.A.1, White discloses methods for independently adjusting supply voltages and clock frequencies in "a multi-core integrated circuit," or "[a] processor or other integrated circuit [that] may include multiple logic cores." EX1006, 2:15-22, 4:31-35, 6:41-46, 8:41-49, 14:39-54, 15:42-64.

[1.1] a first set of processor cores of the multi-core processor,

83. In my opinion, White/Talwar renders obvious Limitation [1.1]. As discussed in Section X.A.4 and shown in modified Figure 2A, below, the proposed White/Talwar combination implements Talwar's teaching that "[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source." EX1007, [0011], Claims 1, 25.

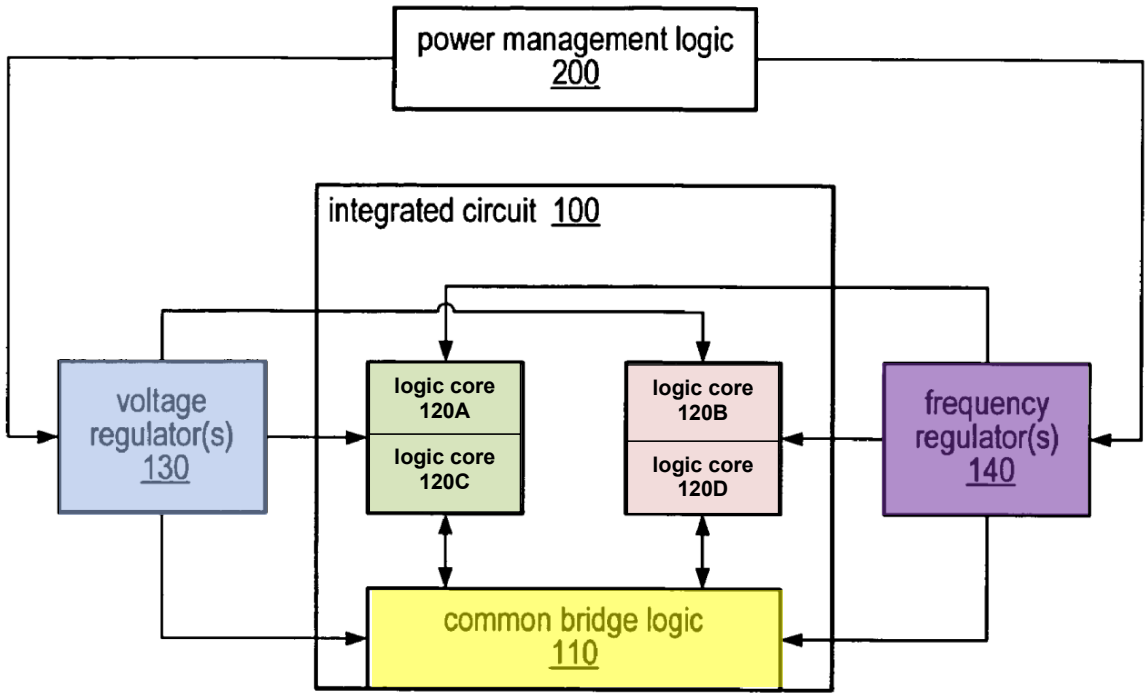

FIG. 2A

84. As shown in modified Figure 2A of White above, cores 120A and 120C of integrated circuit 100 (i.e., the multi-core processor) are grouped into a first cluster that is a first set of processor cores of the multi-core processor and cores 120B and 120D are grouped into a second cluster that is a second set of processor cores of the multi-core processor. *See* EX1006, FIG. 2A (modified).

[1.1.1] wherein each processor core from the first set of processor cores is configured to dynamically receive a first supply voltage and a first output clock signal of a first phase lock loop (PLL) having a first clock signal as input;

85. In my opinion, White/Talwar renders obvious Limitation [1.1.1]. As discussed in Section X.A.1, White teaches that independent supply voltages and

clock signals may be provided to separate processor cores in a multi-core processor. EX1006, 10:11-11:23. White also teaches that the supply voltages and clock signal frequencies should be "adjusted" as needed, for example to support "power management" or "temperature management" requirements. EX1006, 2:15-44. White further teaches that internal clocks for each clock domain may be generated "using a phase-locked loop (PLL) locked to an externally provided clock signal, such as may be provided by frequency regulator 140," which provides "multiple, independent clock signals" to the system. EX1006, 7:40-45, 9:62-65.

86. White discloses that the independent clock signals sent from the frequency regulator to the PLLs at the cores may be generated according to both the First, Second, and Third Configurations discussed above in Section X.A.1. As discussed, those Configurations are:

87. **First Configuration**: Frequency regulator 140 includes "multiple clock sources" in the form of multiple crystal "oscillators" "in order to supply multiple, independent clock signals of different frequencies" (EX1006, 7:40-45), each feeding "a phase-locked loop (PLL)" at one of the clusters.

88. **Second Configuration**: Frequency regulator 140 includes "multiple clock sources" in the form of multiple PLL's, each fed by a corresponding crystal "oscillator" "in order to supply multiple, independent clock signals of different frequencies" (EX1006, 7:40-45), each feeding "a phase-locked loop (PLL)" at one

of the clusters, or directly feeding the cores themselves (*id.*, 9:62-65). EX1008, FIG. 1, Section VII.B.

89. **Third Configuration**: Frequency regulator 140 includes "clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals" (EX1006, 7:40-43), each feeding "a phase-locked loop (PLL)" at one of the clusters, or directly feeding the cores themselves (*id.*, 9:62-65).

90. With respect to the First and Second Configurations, White discloses the use of multiple clock sources within the frequency regulator to generate the "multiple, independent clock signals," where "clock source" is defined to include "an oscillator or phase locked loop (PLL)." EX1006, 7:23-45. Therefore, White discloses multiple oscillators or PLLs within the frequency regulator to generate the "multiple, independent clock signals" provided to the PLLs at each core that generate each core's internal clock. *See id.* As explained above, this means that the "multiple, independent clock signals" originate from two different external crystal oscillators. With respect to the Third Configuration, White also discloses that "clock trees or other circuitry" may be used to "split or divide a clock signal into multiple individual signals." EX1006, 7:40-43. In the Third Configuration, the signals provided to the PLLs at each core are derived from "clock trees or other circuitry configured to split or divide a clock signal into multiple individual signals." *Id.*

Thus, while the clock signals in the Third Configuration share a common source that they both depend on, they are still different signals.

91. Each of the above Configurations satisfy Limitation [1.1.1]. The First and the Second Configurations satisfy this Limitation because they disclose a first input clock signal to a PLL for the first cluster of cores that is independent of a second input clock signal to a PLL for the second cluster of cores, in the sense that there is no dependency between the two input clock signals. The Third Configuration satisfies this Limitation to the extent that "independent" is understood simply to mean different clock signals that can depend from the same reference oscillator.

92. Additionally, as explained in Section X.A.4 and shown in modified Figure 2A, the White/Talwar combination implements White's teachings of dynamically providing supply voltages and clock frequencies with Talwar's clustering of cores, which dynamically receive the same. The proposed combination reduces the amount of on-chip circuitry required to support independent voltage regulation and clock generation for each core, while preserving the power management flexibility provided by White's multiple voltage and clock domains.

93. Thus, the White/Talwar combination discloses this Limitation as follows: each processor core in each cluster "dynamically receive[s]" a common supply voltage ("first supply voltage") and common clock signal ("first output clock

signal"), the common clock signal being derived from a separate PLL for each cluster of cores ("of a first phase lock loop (PLL)"), the PLL being phase locked to the first of the multiple, independent clock signals provided by the frequency regulator in accordance with any of the Configurations discussed above ("having a first clock signal as input").

[1.2] a second set of processor cores of the multi-core processor,

94. In my opinion, White/Talwar renders obvious Limitation [1.2]. As discussed in Limitation [1.1], the White/Talwar combination implements Talwar's teaching that "[t]he cores and voltage sources are divided into clusters, whereby multiple cores in a cluster receive power from a single voltage source." EX1007, [0011], Claims 1, 25. As shown in modified Figure 2A, cores 120B and 120D of integrated circuit 100 (i.e., the multi-core processor) are grouped into a second cluster that is "a second set of processor cores of the multi-core processor."

[1.2.1] wherein each processor core from the second set of processor cores is configured to dynamically receive a second supply voltage and a second output clock signal of a second PLL having a second clock signal as input,

95. In my opinion, White/Talwar renders obvious Limitation [1.2.1]. As discussed in Limitation [1.1.1], the White/Talwar combination includes multiple sets, or clusters of cores, wherein each cluster shares a supply voltage and clock signal that is independent from the supply voltages and clock signals of other clusters. *See* EX1006, 10:11-11:23; EX1007, [0011]-[0012]. Therefore, for the

c) Claim 3

[3.1] The multi-core processor claim 1, wherein the interface block further comprises a second level shifter that is referenced to the first supply voltage and adapted to translate second logic levels associated with the second set of processor cores to first logic levels associated with the first set of processor cores for a second signal traveling from the second set of processor cores to the first set of processor cores.

104. In my opinion, White/Talwar renders obvious Limitation [3.1] for reasons discussed above in Limitation [2.1].

d) Claim 5

[5.1] The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.

105. In my opinion, White/Talwar renders obvious Limitation [5.1]. White discloses that the "logic circuitry or firmware configured to control the operating voltage and frequency of logic cores" may "reside in the processor itself." EX1006, 5:17-24; *see also* 13:7-26 ("logic circuitry…internal…to integrated circuit"), 14:32-38 ("temperature control logic [] and power management logic [] may…be part of a single set of logic."), FIGs. 2A, 2B. Thus, the logic circuitry within the processor may include the "power management logic" that sends control signals to the processor cores by "writ[ing]to one or more control registers" within each processor core. EX1006, 10:7-9. Further, a POSITA would have understood that circuitry design involves layout decisions that balance performance, area, and thermal considerations, and that placement of control logic circuitry along the perimeter of

the multi-core processor or substantially central to processor cores is an obvious design choice. *See also* EX1033, FIG. 8, [0056], [0103]-[0107] (depicting chip "floor plan" with "a controller for signal transmission from/to storage control units SC" that are "located at the side nearer to the SC [storage control units]," which is substantially central to the processor units.). For example, the central placement of shared components, such as control logic, typically shortens the total length of wires between the cores and the control logic, shortening the delay between the cores and the control logic. On the other hand, a POSITA would also understand that it can be more efficient under some circumstances to place the control blocks, such as the power management controller, around the periphery of core clusters, rather than substantially central to the core clusters. That way the core clusters can be placed closed to each other and may have shorter inter-cluster communication wires.

e) Claim 8

[8.1] The multi-core processor of claim 1, wherein the first set of processor cores are located in a first region of the multi-core processor, and the second set of processor cores are located in a second region of the multi-core processor.

106. In my opinion, White/Talwar renders obvious Limitation [8.1]. As discussed in Section X.A.4, White/Talwar includes a first cluster of logic cores 120A and 120C (the first set of logic cores) that share a voltage signal, clock signal, and cluster-level clock control circuitry and a second cluster of logic cores 120B and 120D (the second set of logic cores) that share a different voltage signal, clock

direct control signaling to the cores in the manner taught by Cho, the processor "may optimize both performance as well as thermal control." EX1010, [0047]. Therefore, a POSITA would be motivated to combine White, Talwar, and Cho to gain the predictable benefits disclosed in Cho.

3. Modification of the White/Talwar Combination in View of Cho

127. The functionality of Cho's TCU 102 may be incorporated into White's existing control logic circuitry. White explains that the logic circuitry (i.e., the "temperature control logic 210" and "power management logic 200") "configured to control the operating voltage and frequency of logic cores," "may reside in the processor itself" (EX1006, 5:17-20) and "may both be part of a single set of logic, implemented in hardware, firmware, software, or a combination." *Id.*, 14:32-38.

128. Cho provides additional guidance on the location of the control circuitry beyond White's disclosure of "in the processor" (i.e., integrated circuit 100), illustrating the control circuitry as alternatively substantially central to the functional units (e.g., processor cores) (EX1010, FIG. 1) or in the periphery of the multi-core processor *(id.*, FIG. 4). A POSITA would have understood both of these locations to represent acceptable known design choices. *See* EX1033, FIG. 8, [0056], [0103]-[0107], *see also* EX1036, FIG. 5, 6:39-47.

129. Finally, Cho's "core control signals cctrl1 and ccrtl2" that couple the "TCU" to "processor cores" follow the existing control path signals White describes

to enable "power management logic" to "write to one or more control registers" in the processing cores. EX1006, 10:4-8.

130. Accordingly, as depicted below in the modified portions of Figure 2B, the "control signals cctrl1 and ccrtl2" [red] as taught by Cho could be integrated within the power/temperature management set of logic (i.e., logic circuitry [blue]) in the White/Talwar combination and may be located within the multi-core processor as taught by both White and Cho. *See* EX1010, Fig. 2B (modified below).

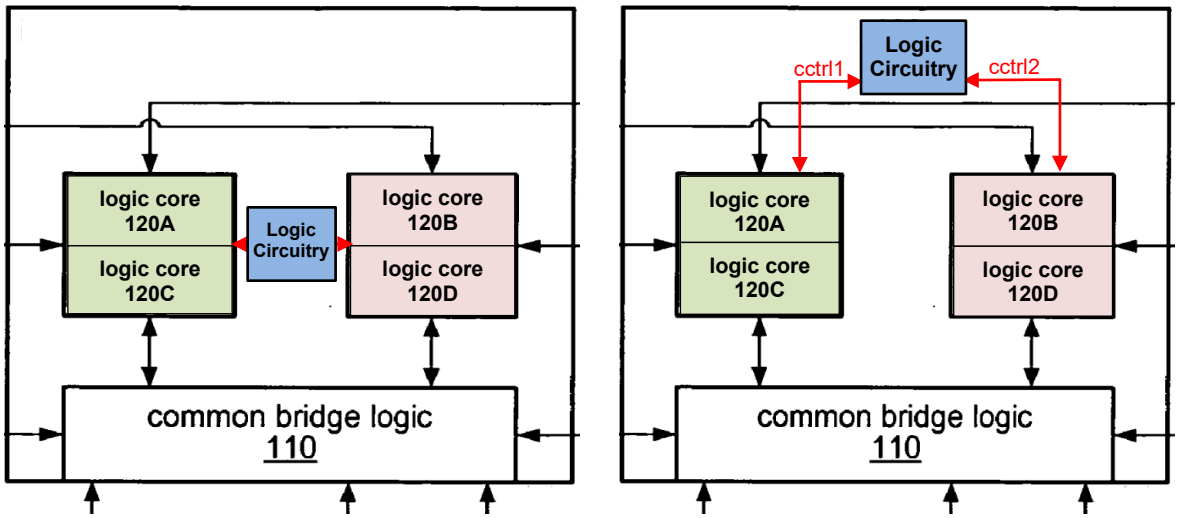

Substantially Central Periphery

EX1006, FIG. 2B (modified). A POSITA would have understood that this single set of logic circuitry [blue] is a control block within the processor that could provide "control signals cctrl1 and ccrtl2," to registers within the processor cores instructing them to re-allocate workloads. The specific implementation of the logic circuitry within White/Talwar combination would not change the operation of White's power

management and temperature control within the multi-core processor. *See* EX1006, Abstract, 5:17-20, 13:7-26, 14:32-38, 16:58-63.

4. Element-by-Element Analysis

a) Claim 5

[5.1] The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a periphery of the multi-core processor.

131. In my opinion, White/Talwar/Cho renders obvious Limitation [5.1]. White discloses that the "logic circuitry or firmware configured to control the operating voltage and frequency of logic cores" may "reside in the processor itself." EX1006, 4:49-5:17-24, 13:7-26. Further, as explained in Section X.B.3, the logic circuitry [blue] provides control signals to the processor cores and is located along the perimeter of the multi-core processor as shown in Fig. 4. *See also* EX1010, Fig. 4, [0038]-[0040], [0042], [0047]-[0048].

b) Claim 14

[14.1] The multi-core processor of claim 1, wherein the first set of processor cores and the second set of processor cores are configured to receive one or more control signals from one or more control blocks located in a common region that is substantially central to the first set of processor cores and the second set of processor cores

132. In my opinion, White/Talwar/Cho renders obvious Limitation [14.1] for reasons discussed in Limitation [5.1]. Further, as explained in Section X.B.3, in addition to locating logic circuitry [blue] along the periphery of a multi-core

statements were made with the knowledge that willful false statement and the like are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Respectfully submitted,

Dated: January 21, 2025

Murali Annavaram, Ph.D.